



11-06-00

11/03/00 THE COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, DC 20231

Attorney Docket: STL0000390JS1/1745P

jc860 U.S. PATENT
11/03/00
jc860 U.S. PATENT
11/03/00
jc860 U.S. PATENT
11/03/00

Sir:

Transmitted herewith for filing is the Patent Application of

Inventor(s): **Jeff W. Josten and Christopher J. Munson**For: **METHOD AND SYSTEM FOR RECOVERING DATA IN A PLURALITY OF SYSTEMS**

Enclosed with the Patent Application are:

- Three (3) sheets of Informal Drawings
- Declaration and Power of Attorney
- Assignment and Recordation Form
- Information Disclosure Statement (PTO Form 1449)
- A certified copy of a _____ application
- Self Addressed, Stamped Postcard

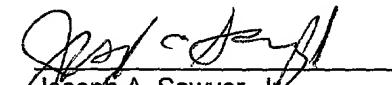
The filing fee has been calculated as shown below:

(Col. 1)(Col. 2)			LARGE ENTITY	
FOR:	NO. FILED	NO. EXTRA	RATE	FEE
BASIC FEE				\$ 710.00
TOTAL CLAIMS	18 - 20	= 0	x 18	= \$ 0.00
INDEP. CLAIMS	3 - 3	= 0	x 80	= \$ 0.00
MULTIPLE DEPENDENT CLAIM PRESENTED			+270	= \$ 0.00
*If the difference in Col. 1 is less than "0", enter "0" in Col. 2			TOTAL	\$ 710.00

- Check no. 2391 in the amount of \$710.00 is enclosed for payment of filing fees. A duplicate of this sheet is attached.

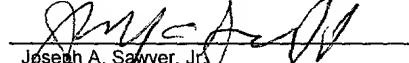
SAWYER LAW GROUP LLP
P.O. Box 51418
Palo Alto, California 94303
(650) 493-4540

Respectfully submitted,


Joseph A. Sawyer, Jr.
Sawyer Law Group LLP
Attorney for Applicants
Reg. No. 30,801

EXPRESS MAIL CERTIFICATE

I hereby certify that the above paper/fee is being deposited with the U.S. Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on November 3, 2000 and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231. "Express Mail" no.: EL547856854US. Signature of Person mailing paper/fee:


Joseph A. Sawyer, Jr.

UNITED STATES PATENT APPLICATION

FOR

**METHOD AND SYSTEM FOR RECOVERING DATA
IN A PLURALITY OF SYSTEMS**

Inventor:

Jeff W. Josten
Christopher J. Munson

Sawyer Law Group LLP
2465 E. Bayshore Road, Suite 406
Palo Alto, California 94303

METHOD AND SYSTEM FOR RECOVERING DATA IN A PLURALITY OF SYSTEMS

FIELD OF INVENTION

The present invention relates generally to computer systems and more specifically to a method and system for recovering data in the event of a system failure.

5

BACKGROUND OF THE INVENTION

Shared File Systems (SFS) is a term applied to IBM's System/390 system for sharing data among virtual machines. IBM's DB2 has been adapted for this type of data sharing in a Multiple Virtual Storage (MVS)/Enterprise Systems Architectures (ESA) environment by using IBM's coupling facility to create multi-system data sharing.

In such a shared system, when one of the systems fails, the update mode locks (data locks) that were held at the time of the failure are "retained" to prevent the other systems from accessing inconsistent data (data that had not yet reached a point of consistency at the time of the failure). To remove the retained data locks, the failed system's logs must be read in a forward and a backward direction in order to bring the data back to a point of consistency. Once this has been done, the retained locks can be removed, and the data is again accessible from all the systems.

One conventional method generally employed to remove the retained locks when an operating system fails is the restart/recovery method. Utilizing the restart/recovery method, the failed system is restarted (either manually or automatically) on another operating system in the cluster and recovery logic is used to "recover" the data being protected by the retained

data locks and bring the data back to consistency. The trouble with this approach is that in order to restart the failed system, a substantial amount of CPU resources could be utilized. Consequently, this use of CPU resources could impose a significant disruption to the work that is already running on the operating system.

5 Accordingly, what is needed is a more efficient method and system for recovering the retained locks of the failed operating system. The method and system should be simple, cost effective and capable of being easily adapted to existing technology. The present invention addresses such a need.

SUMMARY OF THE INVENTION

In a first aspect of the present invention, a method for recovering data in a plurality of systems is disclosed. The method comprises the steps of allowing at least one system of the plurality of systems to fail, retaining a plurality of locks of the at least one system and restarting the at least one system utilizing minimal resources.

In a second aspect of the present invention, a system for recovering data in a plurality of computer systems is disclosed. The system comprises means for allowing at least one computer system of the plurality of computer systems to fail, means for retaining a plurality of locks of the at least one computer system and means for restarting the at least one computer system utilizing minimal resources.

20 According to the present invention, the method and system for recovering retained locks in a plurality of systems recovers the data being protected by the retained locks of a failed system quickly and with minimal system disruption.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is an example of a system in which the present invention could be implemented.

5 Figure 2 is a flowchart of the method in accordance with the present invention.

Figure 3 is a detailed description of step 204 of the flowchart of Figure 2.

DETAILED DESCRIPTION OF THE INVENTION

10 The present invention provides a method and system for recovering data in a plurality of systems. The following description is presented to enable one of ordinary skill in the art to make and use the invention and is provided in the context of a patent application and its requirements. Various modifications to the preferred embodiment will be readily apparent to those skilled in the art and the generic principles herein may be applied to other embodiments. Thus, the present invention is not intended to be limited to the embodiments shown but is to be accorded the widest scope consistent with the principles and features described herein.

15

20 The present invention is disclosed in the context of a preferred embodiment. The preferred embodiment of the present invention provides a method and system for recovering data in a shared data system. In accordance with the present invention, minimal resources are utilized to restart and recover the retained data locks of a failed system. Accordingly, the retained data locks of the failed system are recovered quickly and with minimal system disruption.

For a further description of the present invention, please refer now to Figure 1.

Figure 1 is an example of a system 100 in which the present invention could be implemented. The system 100 comprises a plurality of operating systems 102, 104, 106 wherein each of the plurality of operating systems 102, 104, 106 includes a database management system (DBMS) 103, 105, 107 wherein each of the DBMSs 103, 105, 107 are logically grouped together and operate in tandem with one another. An example of such a system is an IBM S/390 system and DB2 for OS/390 Data Sharing.

In accordance with the present invention a new mode of restarting a failed DBMS is introduced. This new mode (“restart light” mode) preferably specifies that only minimal resources are utilized to perform the restart/recovery process of a failed DBMS. By utilizing minimal resources, the restart/recovery process can be performed quickly and once the data being protected by the retained data locks has been recovered and the data is brought back to consistency, the failed DBMS immediately shuts down in a normal fashion without accepting any new work.

In accordance with the present invention, minimal resources are a predefined plurality of resources that are necessary only for the performance of a restart/recovery process for the failed DBMS. Since the recovery of the data being protected by the retained data locks is the only task that is being performed, any resource that does not facilitate the accomplishment of this task is not needed. For example, a resource that is utilized to enable the failed DBMS to accept new work is not necessary for the performance of the restart/recovery process and is therefore not a minimal resource. The utilization of minimal resources to perform the restart/recovery process serves to significantly reduce the amount of CPU and storage that is required to perform the process and it also reduces the processing

time required to recover the data being protected by the retained data locks. Furthermore, by reducing the CPU and storage requirements, the restart/recovery process can be performed with minimal disruption to the work that is already running on the system.

For a better understanding of the method in accordance with the present invention, please refer now to Figure 2. Figure 2 is a flowchart of the method in accordance with the present invention. In a system comprising a cluster of operating systems wherein each operating system includes a DBMS, the method begins with the abnormal termination (failure) of one of the DBMSs, via step 200. Next, the data locks of the failed DBMS are retained, via step 202. Preferably, the locks are retained by another operating system within the cluster of operating systems. Finally, the failed DBMS is restarted utilizing minimal system resources, via step 204.

For a more detailed understanding of the present invention, please refer now to Figure 3. Figure 3 is a detailed description of step 204 of the flowchart of Figure 2. First, an operating system other than the operating system of the failed DBMS is allowed to restart the failed DBMS, via step 300. Preferably, the operating system restarts the failed DBMS in “restart light” mode after receiving a request to restart the failed DBMS in “restart light” mode. This request is preferably made manually or automatically via computer software. Next, minimal resources of the operating system are utilized to recover the data being protected by the retained locks of the failed DBMS, via step 302. Finally, once the data being protected by the retained locks has been recovered and brought back to consistency, the failed DBMS terminates itself in a normal fashion, via step 304. Preferably, steps 300-304 are performed wherein the failed DBMS does not accept any new work. Once the data

being protected by the retained data locks is released and brought back to consistency, full lock granting protocols are restored throughout the system.

5 Although the preferred embodiment of the present invention is disclosed in the context of being utilized in conjunction with an IBM S/390 system, one of ordinary skill in the art will readily recognize that the present invention could be utilized in conjunction with a variety of systems while remaining within the spirit and scope of the present invention.

10 Such a method may also be implemented, for example, by operating the system 100 to execute a sequence of machine-readable instructions. The instructions may reside in various types of computer readable media. In this respect, another aspect of the present invention concerns a programmed product, comprising computer readable media tangibly embodying a program of machine readable instructions executable by a digital data processor to perform a method for recovering retained locks in a plurality of systems.

15 This computer readable media may comprise, for example, RAM (not shown) contained within the system 100. Alternatively, the instructions may be contained in another computer readable media such as a magnetic data storage diskette and directly or indirectly accessed by the system 100. Whether contained in the system 100 or elsewhere, the instructions may be stored on a variety of machine readable storage media, such as a DASD storage (e.g. a conventional “hard drive” or a RAID array), magnetic tape, electronic read-only memory (e.g., ROM, CD-ROM, EPROM, or EEPROM), an optical storage device (e.g., 20 CD ROM, WORM, DVD, digital optical tape), paper “punch” cards, or other suitable computer readable media including transmission media such as digital, analog, and wireless communication links. In an illustrative embodiment of the invention, the machine-readable instructions may comprise lines of compiled C, C++, or similar language code commonly

used by those skilled in the programming for this type of application arts.

Through the use of the present invention, minimal resources are utilized to perform the restart/recovery process of a failed DBMS. The utilization of minimal resources to perform the restart/recovery process serves to significantly reduce the amount of CPU and storage that is required to perform the process and it also reduces the processing time required to recover the retained data locks. Furthermore, by reducing the CPU and storage requirements, the restart/recovery process can be performed with minimal disruption to the work that is already running on the system.

Although the present invention has been described in accordance with the embodiments shown, one of ordinary skill in the art will readily recognize that there could be variations to the embodiments and those variations would be within the spirit and scope of the present invention. Accordingly, many modifications may be made by one of ordinary skill in the art without departing from the spirit and scope of the appended claims.

CLAIMS

What is claimed is:

1. 1. A method for recovering data in a plurality of systems comprising the steps of:
 2. a) allowing at least one system of the plurality of systems to fail;
 3. b) retaining a plurality of locks of the at least one system; and
 4. c) restarting the at least one system utilizing minimal resources.
1. 2. The method of claim 1 wherein step b) further comprises allowing another system of the plurality of systems to retain the plurality of locks of the at least one system.
1. 3. The method of claim 2 wherein step c) further comprises:
 2. c1) allowing the another system of the plurality of systems to restart the at least one system;
 3. c2) recovering data being protected by the retained locks of the at least one system utilizing minimal resources of the another system; and
 4. c3) allowing the at least one system to terminate in a normal fashion.
1. 4. The method of claim 3 wherein minimal resources consists of a predefined plurality of resources necessary to recover the data being protected by the retained locks of the at least one system.
1. 5. The method of claim 3 wherein step c1) further comprises:

2 c1i) providing a request to restart the at least one system utilizing minimal
3 resources;
4 c1ii) allowing the another system to detect the request;
5 c1iii) allowing the another system to restart the at least one system based on the
6 request.

1 6. The method of claim 1 wherein the plurality of locks comprise a plurality of data
2 locks.

1 7. A system for recovering data in a plurality of computer systems comprising:
2 means for allowing at least one computer system of the plurality of computer systems
3 to fail;
4 means for retaining a plurality of locks of the at least one computer system; and
5 means for restarting the at least one computer system utilizing minimal resources.

1 8. The system of claim 7 wherein the means for retaining the plurality of locks further
2 comprises means for allowing another computer system to retain the plurality of locks of the
3 at least one computer system.

1 9. The system of claim 8 wherein the means for restarting the at least one computer
2 system further comprises:
3 means for allowing the another computer system to restart the at least one computer
4 system;

5 means for recovering data being protected by the retained locks of the at least one
6 computer system utilizing minimal resources of the another computer system; and
7 means for allowing the at least one computer system to terminate in a normal
8 fashion.

1 10. The system of claim 9 wherein minimal resources consists of a predefined plurality
2 of resources necessary to recover the data being protected by the retained locks of the at least
3 one computer system.

1 11. The system of claim 9 wherein means for allowing the another computer system to
2 restart the at least one computer system further comprises:

3 means for providing a request to restart the at least one computer system utilizing
4 minimal resources;

5 means for allowing the another computer system to detect the request;
6 means for allowing the another computer system to restart the at least one computer
7 system based on the request.

1 12. The system of claim 7 wherein the plurality of locks comprise a plurality of data
2 locks.

1 13. A computer readable medium comprising program instruction for recovering data in
2 a plurality of systems, the program instructions comprising the steps of:

3 a) allowing at least one system of the plurality of systems to fail;

- b) retaining a plurality of locks of the at least one system; and
- c) restarting the at least one system utilizing minimal resources.

14. The computer readable medium of claim 13 wherein step b) further comprises allowing another system of the plurality of systems to retain the plurality of locks of the at least one system.

15. The computer readable medium of claim 14 wherein step c) further comprises:

- c1) allowing the another system of the plurality of systems to restart the at least one system;
- c2) recovering data being protected by the retained locks of the at least one system utilizing minimal resources of the another system; and
- c3) allowing the another system to terminate the at least one system in a normal fashion.

16. The computer readable medium of claim 15 wherein minimal resources consists of a predefined plurality of resources necessary to recover the data being protected by the retained locks of the at least one system.

17. The computer readable medium of claim 15 wherein step c1) further comprises:

- c1i) providing a request to restart the at least one system utilizing minimal resources;
- c1ii) allowing the another system to detect the request;

c1iii) allowing the another system to restart the at least one system based on the request.

18. The computer readable medium of claim 13 wherein the plurality of locks comprise a plurality of data locks.

ABSTRACT

In a first aspect of the present invention, a method for recovering data in a plurality of systems is disclosed. The method comprises the steps of allowing at least one system of the plurality of systems to fail, retaining a plurality of locks of the at least one system and restarting the at least one system utilizing minimal resources. In a second aspect of the present invention, a system for recovering data in a plurality of computer systems is disclosed. The system comprises means for allowing at least one computer system of the plurality of computer systems to fail, means for retaining a plurality of locks of the at least one computer system and means for restarting the at least one computer system utilizing minimal resources. According to the present invention, the method and system for recovering retained locks in a plurality of systems recovers the data being protected by the retained locks of a failed system quickly and with minimal system disruption.

10

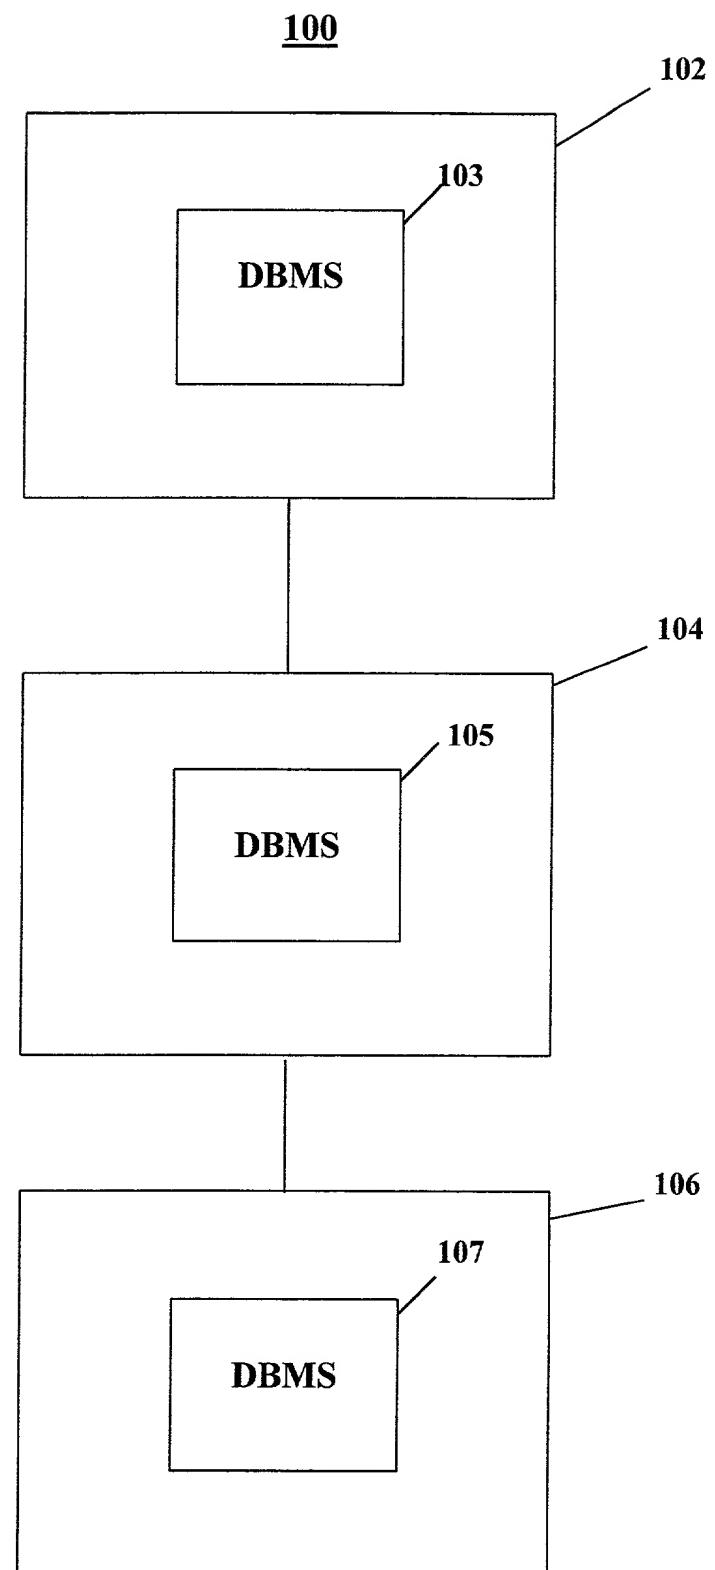


Figure 1

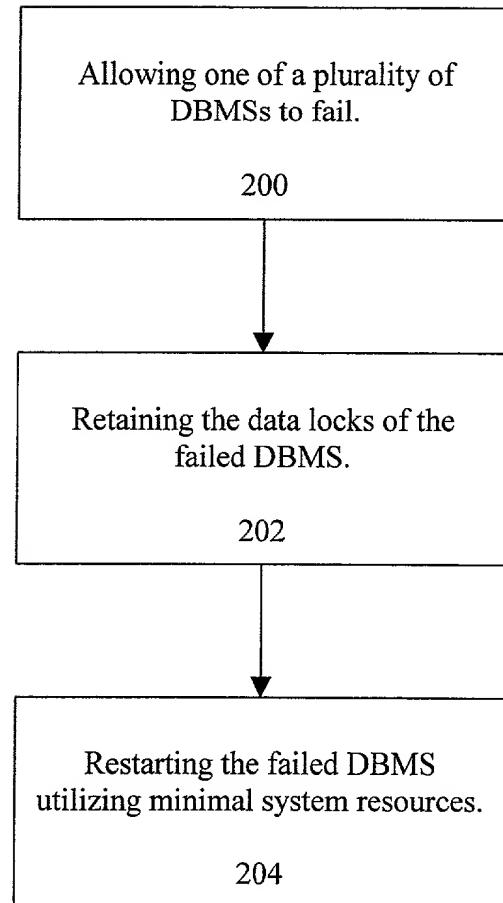


Figure 2

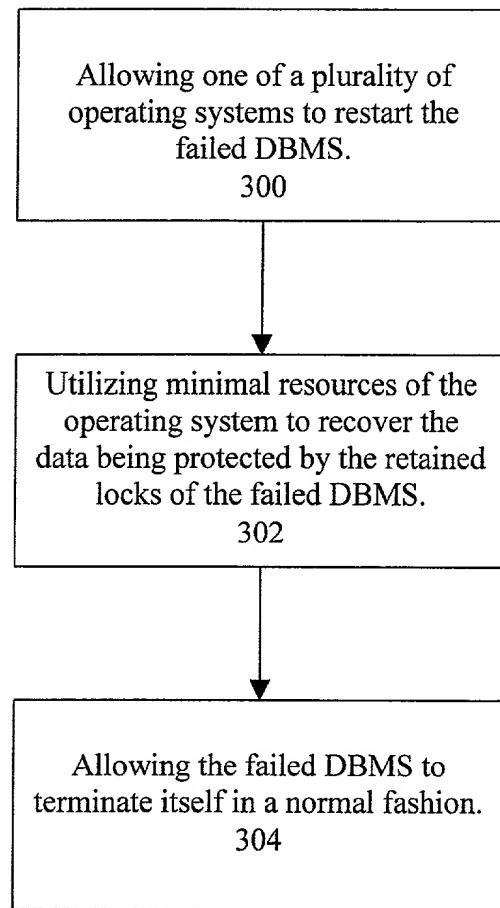


Figure 3

DECLARATION AND POWER OF ATTORNEY FOR UTILITY PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name,

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD AND SYSTEM FOR RECOVERING DATA IN A PLURALITY OF SYSTEMS

the specification of which

X is attached hereto.
 was filed on

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above. I do not know and do not believe that the same was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and said invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months prior to this application.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, Section 1.56 (a).

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119, of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)			Priority Claimed	
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial No.)	(Filing Date)	(Status – patented, pending, abandoned)
(Application Serial No.)	(Filing Date)	(Status -- patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith (list name and registration number).

Romualdas Strimiatis, Reg. No. 35,697
Timothy M. Farrell, Reg. No. 37,321
Ingrid M. Foerster, Reg. No. 36,511
Prentiss W. Johnson, Reg. No. 33,123
Christopher A. Hughes, Reg. No. 26,914
John E. Hoel, Reg. No. 26,279
Edward A. Pennington, Reg. No. 32,588

Joseph C. Redmond, Jr., Reg. No. 18,753
Joseph A. Sawyer, Jr., Reg. No. 30,801
Janyce R. Mitchell, Reg. No. 40,095
Stephen G. Sullivan, Reg. No. 38,329
Michele Liu, Reg. No. 44,875
Wendell J. Jones, Reg. No. 45,961
Doretha L. Robinson, Reg. No. 45,048

as my attorney with full power of substitution and revocation to transact all business in the United States Patent and Trademark Office connected herewith.

Please direct all telephone calls to Mr. Joseph A. Sawyer, Jr. at telephone number (650) 493-4540 and address all correspondence to: **SAWYER LAW GROUP LLP, P.O. Box 51418, Palo Alto, California 94303**

Full Name of **first**/joint Inventor: **Jeff W. Josten**
Residence Address: **6728 Crystal Springs Drive, San Jose, California 95120**
Country of Citizenship: **United States of America**
Post Office Address: **Same**

Signature of Inventor:



Date: 11/1/2000

Full Name of **second**/joint Inventor: **Christopher J. Munson**

Residence Address: **236 Gregg Drive, Los Gatos, California 95032**
Country of Citizenship: **United States of America**
Post Office Address: **Same**

Signature of Inventor:



Date: 11/2/2000
